

REMARKS

Claims 1-69 are pending. Claims 70 to 72 have been newly added. Claims 1-69 including independent claims 1, 22, 44, 63, and 67 are rejected under 35 U.S.C. 102(e) as being anticipated by Fawaz (USP 6,970,424). Independent claims 1, 63, and 67 and dependent claim 10 have been amended to facilitate prosecution. Claims 18, 22-62 have been canceled to facilitate prosecution. Claims 70-72 are new.

The amendments to the independent claims and the new claims are believed to introduce no new matter and are believed supported at least by Figures 4, 6, 7, and 8 and associated description. For example, “Even more accurate congestion control can be implemented by using different quench messages or by providing information in each quench message. Figure 6 is a graphical representation depicting the generation of various quench messages at a congested switch based on buffer or queue levels. When the buffer level exceeds a high threshold 607, a path quench can be generated. A path quench can instruct all switches between the source node and the congested switch to immediately stop sending traffic associated with a particular flow towards the congested switch. A path quench can be used when the buffer level at the congested switch is near capacity and an immediate traffic reduction is necessary to avoid the depletion of credits and the buffer-to-buffer credit mechanism. A high threshold can vary based on buffer size and the importance of avoiding the depletion of credits. In one example, the high threshold is set at 90 percent of the buffer size.” (page 15, line 29 – page 16, line 7)

“If the buffer level is between a high threshold 607 and a low threshold 603, an edge quench message can be generated by the congested switch. The edge quench instructs the congestion control switch closest to the source node to reduce the traffic flow towards the congested switch. Network nodes that can recognize instructions from another network node to reduce traffic flow are herein referred to as congestion control switches. The edge quench messages can provide a gradual reduction in traffic flow from a particular source node. A cut off threshold 601 can also be provided to prevent the generation of quench messages. By not sending a quench message, congestion control and slow down of other traffic can be avoided when the quench message is not necessary.” (page 16, line 8 – page 16, line 18)

Independent claims 1, 63, and 67 now recite “wherein the first instruction is sent only to the first intermediate switch to reduce transmissions at the first intermediate switch and the

second instruction is sent to the plurality of switches including the first intermediate switch to reduce transmissions at the plurality of switches including the first intermediate switch.” None of the references including Fawaz cited by the Examiner is believed to teach or suggest this recitation. For example, Fawaz does not teach or suggest generating two different types of instructions based on the amount of congestion control needed. Furthermore, Fawaz does not teach or suggest sending one type of instruction to only an edge switch while the other type of instruction is sent to a plurality of switches including the edge switch.

Fawaz describes “Service Level Agreements, or SLA's, are defined between pairs of Packet Switches and guarantee a minimum quality-of-service (minimum bandwidth) between the two packet switches. When a packet arrives at a node from a packet switch, the node inspects certain classification information contained within the packet. In one embodiment, such classification information is the source and destination identifiers (e.g., addresses) of the packet, while in other embodiments classification information additionally includes other information. Using the classification information, the packet classifies the packet with an SLA. A scheduler in the node ensures that packets from each SLA are scheduled for transmission at at least the minimum data rate corresponding to the SLA.” (column 4, lines 23-37)

Fawaz only describes a particular node characterizing traffic and the same node using a “scheduler in the node” to ensure that service level agreements are met (SLA). Fawaz controls traffic at a scheduler at the same node where the traffic characterization takes places. Although it is theoretically possible that Fawaz uses a switch and two types of instructions with each type of instruction sent to different entities to control congestion, this is not taught or suggested and can not be assumed. Fawaz seems to suggest that congestion control occurs at the scheduler in the node. It is acknowledge that there are congestion control mechanisms available, however, these congestion control mechanisms such as that described in Fawaz all lack elements recited in the independent claims.

New claim 70 recites “a processor operable to characterize traffic flow and determine buffers levels at the interface, wherein buffer levels exceeding a high threshold triggers the generation of a path quench frame that is sent to the plurality of fibre channel switches including the edge switch to limit traffic flow to the interface and wherein buffer levels between a low threshold and a high threshold triggers the generation of an edge quench frame that is send to the edge switch to reduce traffic flow to the interface from the edge switch.”

Fawaz does not teach or suggest determining buffer levels and does not teach or suggest sending an edge quench frame based on one buffer level and a path quench frame based on another buffer level. Furthermore, Fawaz does not use two types of instructions with each type of instruction sent to different entities to control congestion. It is again theroretically possible that Fawaz examines buffer levels even though it is not suggested, however, this can not be assumed. The SLA agreements described in Fawaz can be enforced without any determination of buffers levels, as SLA agreements can be driven simply by monitoring the transmission rates of data.

Furthermore, Fawaz does not teach or suggest any fibre channel switch as recited in the claims. The Examiner argues that Fawaz recites “Further, the actual links between the QoS Nodes can be formed in any manner known to those of skill in the art. For instance, the links interconnecting the QoS Nodes can be built from single or multiple twisted wire pairs, optical fibers, or coaxial cable. In other embodiments, the links can be radio links, or even free space infrared links. In addition the protocol used by the links may be based on Gigabit Ethernet, ATM, WDM, SONET, or other technology.” However, Fawaz never specifically describes any fibre channel switch and only makes passing reference to other technologies.

Fibre channel switches operate in a different manner than conventional Ethernet switches. Technologies can not be readily cross applied between the different switch types. For example, fibre channel switches do not allow for out of order delivery or ready packet dropping as is allowed by Ethernet switches and routers. As noted in the specification, “Many conventional network protocols use packet dropping to alleviate congestion at a network node. In one example, a network node in an IP based network receives input data from multiple sources at a rate exceeding its output bandwidth. In conventional implementations, selected packets are dropped to allow transmission of remaining packets within the allocated output bandwidth. Packets can be dropped randomly or by using various selection criteria. The dropped packets are ultimately retransmitted under the control of a higher level protocol such as TCP.

In networks such as fibre channel networks, packet dropping is generally not allowed. Instead, networks such as fibre channel networks implement end-to-end and buffer-to-buffer flow control mechanisms. End-to-end and buffer-to-buffer flow control mechanisms do not allow a first network node to transmit to a second network node until a second network node is ready to receive a frame. The second network node typically indicates that it is ready to receive

a frame by granting credits to the first network node. When frames are transmitted, credits are used. When no credits remain, the first network node can no longer transmit to the second network node. However, end-to-end and buffer-to-buffer flow control mechanisms provide only a very rough technique for controlling congestion, as the mechanism blocks all traffic along a particular link. Such blocking can quickly propagate upstream to other links in a fibre channel network topology. Some of these links might serve as corridors for paths that do not include the originally congested link. Hence, congestion at one link of one network path can sometimes cause blocking over a much wider portion of a fibre channel topology.” (page 1, line 24 – page 2, line 4)

Fawaz in fact alludes to difficulties in applying the technology to other types of networks. “Yet, the frame structure for Ethernet (or other) packets does not need to be modified as they would for SONET and ATM, causing the system to appear to the packet switch as a shared Ethernet. Nor does a system in accordance with the invention require complex hardware and software akin to that required for ATM.” (column 10, lines 56-62) Fawaz does not even begin to address or even recognize difficulties of applying the Fawaz technology to fibre channel switches.

In light of the above remarks, the rejections to the independent claims are believed overcome for at least the reasons noted above. Applicants’ Representative believes that all pending claims are allowable in their present form. If the Examiner has any questions or concerns for Applicants’ Representative, the Examiner is encouraged to contact her at the number provided below.

Respectfully submitted,
BEYER WEAVER LLP

/ G. Audrey Kwan/
Godfrey Audrey Kwan
Reg. No. 46,850

Beyer Weaver, LLP
P.O. Box 70250
Oakland, CA 94612-0250
(510) 663-1100